

Concurrent Cube-and-Conquer[★]

Peter van der Tak¹, Marijn J.H. Heule^{1,2}, and Armin Biere³

¹ Delft University of Technology, The Netherlands

² University of Texas at Austin, United States

³ Johannes Kepler University Linz, Austria

Abstract. Recent work introduced the cube-and-conquer technique to solve hard SAT instances. It partitions the search space into cubes using a lookahead solver. Each cube is tackled by a conflict-driven clause learning (CDCL) solver. Crucial for strong performance is the cutoff heuristic that decides when to switch from lookahead to CDCL. Yet, this offline heuristic is far from ideal. In this paper, we present a novel hybrid solver that applies the cube and conquer steps simultaneously. A lookahead and a CDCL solver work together on each cube, while communication is restricted to synchronization. Our concurrent cube-and-conquer solver can solve many instances faster than pure lookahead, pure CDCL and offline cube-and-conquer, and can abort early in favor of a pure CDCL search if an instance is not suitable for cube-and-conquer techniques.

1 Introduction

Current satisfiability solvers that target industrial instances are almost always based on the conflict-driven clause learning (CDCL) [7] technique. A technique with fast heuristics and data structures that can successfully solve very large instances by propagating decisions and learning additional information when conflicts arise. Yet on small, hard problems lookahead solvers [3] perform better by applying much more reasoning in each search node and then recursively splitting the search space until a solution is found.

Recent work [4] has shown that the two techniques can be combined successfully, resulting in better performance particularly for very hard instances. The key insight is that lookahead solvers can be used to partition the search space into subproblems that are easy for a CDCL solver to solve. By first partitioning (*cube*) and then solving each subproblem (*conquer*), some instances can be solved within hours rather than days. This cube-and-conquer approach, particularly the *conquer* part, is also easy to parallelize.

The challenge to make this technique work lies in developing good heuristics to determine when to stop partitioning and start solving. The current heuristics already give good results, but are far from optimal and require some fine tuning to work well with instances of different difficulty. For example, applying too much partitioning can undesirably increase the run time of otherwise easy instances.

[★] The second author is supported by DARPA contract number N66001-10-2-4087. The third author is supported by FWF, NFN Grant S11408-N23 (RiSE).

The most important problem in developing a better heuristic is that in the partitioning phase no information is available about how well the CDCL solver will perform on a subproblem. The heuristic is required to estimate this performance, but this is not always reliable. In this work we use an online approach that runs both phases concurrently, and that thereby avoids this problem. We focus less on the parallelization of the *conquer* phase.

Other than improving the performance of cube-and-conquer by replacing this heuristic, the online approach aims to solve another problem: for some instances cube-and-conquer performs worse than CDCL regardless of the configuration of the solvers and heuristics. Our approach is able to quickly identify these instances, in which case the problem can be solved using a classical CDCL search.

We believe that CCC is particularly interesting as part of a portfolio solver, where our predictor can be used to predict whether to apply cube-and-conquer techniques. The authors of SATzilla specifically mention in their conclusion that identifying solvers that are only competitive for certain kinds of instances still has the potential to further improve SATzilla’s performance substantially [10].

2 Preliminaries

For a Boolean variable x , there are two *literals*, the positive literal, denoted by x , and the negative literal, denoted by $\neg x$. A *clause* is a disjunction of literals, and a *CNF formula* is a conjunction of clauses. A clause can be seen as a finite set of literals, and a CNF formula as a finite set of clauses. A *unit clause* contains exactly one literal. A truth assignment for a CNF formula F is a function φ that maps variables in F to $\{\mathbf{t}, \mathbf{f}\}$. If $\varphi(x) = v$, then $\varphi(\neg x) = \neg v$, where $\neg \mathbf{t} = \mathbf{f}$ and $\neg \mathbf{f} = \mathbf{t}$. A clause C is satisfied by φ if $\varphi(l) = \mathbf{t}$ for some $l \in C$. An assignment φ satisfies F if it satisfies every clause in F . A *cube* is a conjunction of literals and a *DNF formula* a disjunction of cubes. A cube can be seen as a finite set of literals and a DNF formula as a finite set of cubes. If $c = (l_1 \wedge \dots \wedge l_k)$ is a cube, then $\neg c = (\neg l_1 \vee \dots \vee \neg l_k)$ is a clause. A truth assignment φ can be seen as the cube of literals l for which $\varphi(l) = \mathbf{t}$. A cube c is satisfied by φ if $\varphi(l) = \mathbf{t}$ for all $l \in c$. An assignment φ satisfies DNF formula D if it satisfies some cube in D . A DNF formula D is called a *tautology* if every full assignment φ satisfies D .

2.1 Cube-and-conquer

The technique proposed in this work is based on *cube-and-conquer* (CC) [4]. CC was designed for solving very hard instances by partitioning the search space into *cubes* using a lookahead solver (`march_cc`), and then solving each cube using an incremental CDCL solver (iLingeling). The key observation made by the authors is that CDCL solvers often solve these cubes very fast, and as a result the two-phase solver is faster than either solver on its own. Additionally, it is natural to parallelize by solving multiple cubes in parallel. In this work we mainly use MiniSAT 2.2 [2] as CDCL solver instead of Lingeling, since it is easier to extend.

Cube-and-conquer modifies the lookahead solver to cut off its search based on a *cutoff heuristic*. When the heuristic triggers, the conjunction of decision literals (cube) is stored and the solver continues as if the branch was unsatisfiable. When finished, all cubes are solved incrementally by a CDCL solver, by adding a cube’s literals as assumptions [1] to the original formula and running the search. The disjunction of cubes is a tautology, so that solving each cube individually is equivalent to finding a solution to the original formula.

The cutoff heuristic multiplies the number of assigned decision variables and the total number of assigned variables as an indication of the complexity of the current cube. If this number exceeds a threshold value, the branch is cut off. The threshold is chosen dynamically: it is decreased when lookahead proves unsatisfiability for a branch, because CDCL would likely have solved it faster, and it is also decreased when lookahead descends too deep in the decision tree, which would result in too many cubes. A more detailed explanation follows in Sec. 4.3.

3 Motivation

Cube-and-conquer shows strong performance on several hard application benchmarks [4], beating both the lookahead and CDCL solvers that were used for the cube and conquer steps. However, on many other instances, either lookahead or CDCL outperforms CC. We observed that for benchmarks for which CC has relatively weak performance, two important assumptions regarding the foundations of CC do not hold in general.

First, in order for CC to perform well, lookahead heuristics must be able to split the search space into cubes that, combined, take less time for CDCL to solve. Otherwise, cube-and-conquer techniques are ineffective and CDCL would be the preferred solving technique. Second, lookahead must be able to refute cubes that are easy for CDCL to solve, and it should not refute cubes that are still hard for CDCL. When this assumption fails, the cutoff heuristic will perform badly, and the cube phase either generates too few cubes and leaves a potential performance gain unused, or generates too many cubes because cubes with fewer decisions are also easy for CDCL to solve. In this section, we discuss the involved heuristics in more detail, and we discuss how to predict when these heuristics are ineffective.

In related work on portfolio SAT solving [10,9,6] machine learning techniques are used for selection (including parameters) and scheduling of SAT solvers. These techniques are based on measuring several features of instances, which can be characterized as either being static, such as number of variables and clauses, or dynamic, such as the number of propagated assignments at certain decision depths (local search or DPLL probing [10]). In this section, we describe a new dynamic feature, which allows us to predict the effectiveness of lookahead and cutoff heuristics in the context of CC and extensions.

3.1 Lookahead heuristics

To compare the performance of CDCL and CC, we ran both solver types⁴ on all application benchmarks of SAT 2009. CDCL was able to solve 57 more benchmarks than CC within the timeout of 900 seconds (171 vs 114). For some instances, the performance gap was huge (in favor of CDCL), in particular on satisfiable ones. This can be explained as follows. After a decision, the reduced formula might be harder (or at least not easier) than the original one. This may be caused by ineffective lookahead heuristics. In case a decision hardly reduces the search space, the conquer solver could need to solve two similar problems instead of one, thereby raising the computational costs. On satisfiable formulas this negative effect is expected to be larger, since a single wrong decision might bring the solver in a part of the search space without solutions.

The main reason for this negative effect is that the key assumption underlying CC fails. This assumption expects that lookahead decision heuristics can select for a formula F a decision variable x in such a way that $F \cup \{x\}$ and $F \cup \{\neg x\}$ are easier to solve separately than F itself. It was shown that for several benchmarks this assumption holds [4]. However, the results above show that for many benchmarks in the SAT 2009 application suite this is not the case. For those, one would like to apply pure CDCL instead of CC.

Ineffective lookahead heuristics can be observed as follows. Given a formula F and a decision variable x , lookahead creates two branches $F \cup \{x\}$ and $F \cup \{\neg x\}$. The branch that reduces the formula the most is called the *right* branch, or a *discrepancy*. In case lookahead heuristics are effective, then with each decision, but especially each discrepancy the formula becomes much simpler. Thus, after only a few discrepancies, lookahead (or CDCL) should be able to refute the branch. A cube that is reached through many discrepancies suggests that the lookahead heuristics have not been effective for that branch.

3.2 The cutoff heuristic

The cutoff heuristic is crucial for performance of cube-and-conquer. Cutting off too early wastes a potential performance gain, but cutting off too deep can result in a large number of instances increasing the total run time. Yet with the current heuristic it is often the case that thousands or millions of cubes are solved almost instantly, while one or two remain and take the majority of the run time. This suggests that the heuristic is not able to properly detect which branches are easy and should be cut off.

In case this behavior is observed, two complementary actions would be preferred. On the one hand, for the many cubes that are solved almost instantly, the cutoff should have taken place earlier (at a smaller cube) to reduce the cost of the cube phase of CC. On the other hand, for the cubes that require lots

⁴ MiniSAT 2.2 for CDCL; MiniSAT 2.2 and march_cc (cube phase) and iMiniSAT 2.2 (conquer phase) for CC. All benchmarks were first preprocessed using Lingeling as suggested for CC in [4]. We used the same version of Lingeling as in [4].

of computational resources, the cutoff should have been performed later to use lookahead for further partitioning. In short, if this happens – in CC only in the conquer phase – then the cutoff heuristic should be considered ineffective.

3.3 Predicting when to apply cube-and-conquer

To predict for which benchmarks CC is competitive, we propose *concurrent cube-and-conquer* (CCC) as follows. During the cube phase of CC, run a CDCL solver in parallel which follows the decisions of the cube solver (details are described in Sec. 4). By running both solvers simultaneously, the cutoff heuristic becomes obsolete, because the CDCL solver naturally determines whether a cube is easy for CDCL⁵. With the cutoff heuristic out, we only need to predict when lookahead heuristics are ineffective. The following two metrics can be used to predict when this is the case.

First, lookahead techniques appear effective if they can solve some cubes faster than CDCL. While running the lookahead and CDCL solver in parallel, we count the number of times that lookahead is faster than CDCL. For benchmarks for which this count is increased very slowly, say less than once per second, we observed that CC was generally not an effective solving strategy.

Second, if the variable heuristics are effective then each discrepancy should result in a large reduction of the formula. Hence after a certain number of discrepancies the solver should be able to refute that branch. Preliminary experiments suggest that if CCC finds a leaf with over 20 discrepancies early in the search-tree, then lookahead variable heuristics should be considered as ineffective. Lookahead solvers solve the left branch first as it is heuristically most likely to be satisfiable. In contrast, CCC considers the right branch first so that it can quickly detect if a branch with a large number of discrepancies is encountered.

To predict whether an instance is suitable for (C)CC, these metrics are combined as follows. Run CCC and abort it if it enters a branch with more than 20 discrepancies. If after 5 seconds CCC is still running but 10 or fewer cubes were solved by lookahead, also abort the solver. For aborted instances (unpredicted instances), a pure CDCL search is run instead. For instances that were not aborted (predicted instances) CCC is the preferred solving technique and can continue. The same instances usually work well for CC, but they cannot be detected as easily because CDCL is only used in the conquer phase. In fact, we have not been able to come up with a quick CC-based predictor that works well.

4 Concurrent cube-and-conquer

This section describes the *concurrent cube-and-conquer* (CCC) technique. We first describe CCC_∞, and extend it later by adding a cutoff heuristic like in CC for better resource utilization. CCC_∞ constructs a decision tree via the lookahead

⁵ Still, cutoff heuristics can lead to reduced resource usage and better performance, as described in Sec. 4.3.

solver and simultaneously runs a CDCL solver on the newest node of this decision tree. Whenever the lookahead solver assigns a decision variable, the new literal is sent to the CDCL solver, which adds it as an assumption and restarts. This is repeated recursively until either solver proves unsatisfiability, which means that the cube is refuted and both solvers backtrack. Whereas CC uses a cutoff heuristic to determine which branches are cut off, CCC_∞ cuts branches off implicitly when CDCL proves unsatisfiability before lookahead makes another decision.

Ideally, this approach is implemented within one solver. However, due to lack of appropriate data structures, current CDCL solvers only apply lookahead and other forms of preprocessing at the top-level, and not under assumptions. For instance tree-based lookahead [5] requires access to all binary clauses at all decision levels, which can only be accessed in a fast manner by either using full occurrence lists or three watches for non-binary clauses. Both techniques are not easy to combine with data structures currently used in CDCL solvers.

On the other hand, lookahead solvers lack data structures for conflict analysis and learning, which is essential in CDCL solvers for allowing non-chronological backtracking and for cutting off repeated parts of the search. CC and CCC can be seen as two different ways of solving this dilemma by running both types of solvers separately, sequentially in CC and concurrently in CCC_∞ .

CC showed to be particularly useful if many cubes were generated, which means that CCC needs frequent synchronization. To keep the synchronization costs small, CCC_∞ uses asynchronous message queues, where both solvers are peers. This architecture also makes it easy to integrate other solvers in the future.

The solvers in CCC_∞ communicate using two queues: the decision queue Q_{decision} and the result queue Q_{solved} . Whenever the lookahead solver assigns a decision variable, it pushes the tuple $\langle \text{cube } c_{id}, \text{literal } l_{\text{dec}}, \text{backtrackLevel} \rangle$ comprising a uniquely allocated id , the decision literal, and the number of previously assigned decision variables (backtrackLevel). When the CDCL solver reads the new decision from the queue, it already knows all previous decision literals, and only needs to backtrack to the backtrackLevel and add l_{dec} as an assumption to start solving c_{id} . The id is used to identify the newly created cube.

If the CDCL solver proves unsatisfiability of a cube before it receives another decision, it pushes the c_{id} of the refuted cube to Q_{solved} . The solver then continues with the parent cube, by backtracking to the level where all but the last decision literal were assigned. When the lookahead solver reads the c_{id} from Q_{solved} , it backtracks to the level just above this cube's last decision variable and continues its search as if it proved unsatisfiability of the cube by itself.

The CDCL solver proves unsatisfiability of a cube if it encounters a complementary assignment when attempting to assign one of a cube's literals. This is not necessarily the last literal of the cube, so that it may refute not only the cube corresponding to the latest decision read from Q_{decision} , but also one or more of its parent cubes. Therefore, it sends only c_{id} of the smallest cube which it refuted, which implies that the sub cubes are also unsatisfiable.

To keep track of the cubes that are pending to be solved, both solvers keep the trail of decision literals (or assumptions for the CDCL solver) and the ids of the

cubes up to and including each decision literal (or assumption). Whenever either solver proves unsatisfiability of the empty cube, or when it finds a satisfying assignment, the other solver is aborted.

It is possible that the lookahead solver already proved unsatisfiability of a cube when it receives the same result from the CDCL solver. The id is used to discard results on Q_{solved} for cubes that have already been closed. Similarly, it is possible that the lookahead solver makes a decision even though the CDCL solver already proved unsatisfiability of a parent of that cube. In that case the CDCL solver can discard the obsolete item on Q_{decision} .

4.1 Example

Consider the decision tree in Figure 1. The decisions made by the lookahead solver are displayed on the edges, and each node contains the c_{id} of the cube corresponding to the literals on the path from the root of the tree up to that node. The id 's are incremented based on depth first search.




Fig. 1. Example decision tree.

Assume that c_4 has been refuted previously, and both solvers are currently solving c_6 . Now, if the CDCL solver finds a conflict when assigning assumption $\neg x_3$, it knows that c_3 is unsatisfiable and pushes c_3 to Q_{solved} . It then removes assumptions $\neg x_4$, $\neg x_7$, and $\neg x_3$, continues with c_2 , and propagates x_3 because it analyzed the conflict and learned something like $(\neg x_2 \vee x_3)$.

When the lookahead solver reads c_3 from Q_{solved} , it will abort its search in c_6 , skip over c_7 and also abort c_5 , and c_3 . It continues solving c_8 by making decision x_3 from c_2 . Note that it is possible that when the lookahead solver reads c_3 from Q_{solved} , it has already progressed and is solving c_7 , or even c_8 or c_9 . When solving c_7 , the same action can be taken: abort c_7 , c_5 , and c_3 . In case

it is at c_8 or beyond, then c_3 will no longer be part of the trail and the message is skipped because c_3 is already known to be unsatisfiable.

Now consider what would happen if the lookahead solver proves unsatisfiability of c_6 : without sending anything to the CDCL solver, it would backtrack to c_5 and then enter the right branch, pushing $(c_7, x_4, 3)$ onto Q_{decision} . If the CDCL solver has not yet solved c_5 by the time it reads from Q_{decision} , it backtracks to level 3 (c_5), decides x_4 , and thereby starts solving c_7 .

4.2 Implementation

The listing in Fig. 2 shows pseudocode for the implementation of the lookahead solver in CCC_∞ . The function is called recursively for each cube that is entered, and all arguments except for the formula F are initially empty lists. Line 1 allocates a new, unique *id* for the cube. Line 2 checks if the CDCL solver has proved any new cubes unsatisfiable. If it did, and if that cube is part of the cube that is currently being solved, the search for this cube is aborted on line 3. If it is not a parent of the current cube, then this result is no longer relevant and is removed from the queue via line 4.

If the cube was not yet solved by CDCL, line 5 adds its id to the id trail S_{id} , which is the list of all nodes on the path from the root of the decision tree to the current cube. Line 7 sends the id and new decision literal to the CDCL solver, as well as the level at which the literal was added, because the CDCL solver should backtrack to this level before adding the literal. Line 6 only handles the special case of the root cube, which has no decision literal and does not need to be sent to the CDCL solver. Lines 8 to 13 describe the core of any typical lookahead solver, and are responsible for propagating unit clauses, checking if a solution is found, and otherwise making a new decision and calling *LA_search* recursively for the two cubes that result by adding the literal or its complement to φ_{dec} .

Fig. 3 lists the pseudocode for CCC_∞ 's CDCL solver. S contains the ids of cubes that are currently being solved, much like S_{id} in the lookahead solver, and is initially empty. Lines 4 to 9 handle decision literals sent by the lookahead solver. If a new decision is available on queue Q_{decision} , then the solver pushes the id of the new cube on S , backtracks to the indicated level, and adds the new decision literal on lines 7, 8, and 9 respectively. Line 6 handles a special case that can occur when the lookahead solver makes a decision while the CDCL solver already proved its parent cube unsatisfiable. In that case, *backtrackLevel* (the size of the parent cube) will be larger than the number of decisions in the current cube, $|S|$, which indicates that the decision is no longer relevant and can be ignored. If no decisions are waiting on the queue, line 11 makes a decision using the CDCL solver's heuristics. Line 16 will detect if the CDCL solver finds a conflict while assigning one of the literals in the cube, in which case lines 17 and 18 will notify the lookahead solver. Since one conflict can prove multiple cubes in S unsatisfiable, line 17 removes the larger cubes so that line 18 only sends the smallest cube that was proved unsatisfiable. The lookahead solver only needs the smallest cube, because this will implicitly abort the larger cubes too.

The remaining lines are like any other CDCL solver. Restarts are not part of the pseudocode, but should be implemented by backtracking to level $|S|$ instead of level 0 for CCC_∞ to work correctly. In addition, we reset the restart strategy and reduce the clause database size every time a cube is refuted for better performance. Cubes are solved in the same order as they were generated, and two threads of a parallel solver never solve the same cube (the multijob strategy [4]).

```

global  $Q_{\text{decision}} = \emptyset, Q_{\text{solved}} = \emptyset, id = 0$ 

 $LA\_search(F, \varphi_{\text{dec}}, \varphi_{\text{imp}}, S_{\text{id}})$ 
1    $id := id + 1$ 
2   while not  $Q_{\text{solved}}.\text{empty}()$  do
3       if  $Q_{\text{solved}}.\text{head}() \in S_{\text{id}}$  then return UNSAT
4        $Q_{\text{solved}}.\text{remove}()$ 
5        $S_{\text{id}}.\text{push}(id)$ 
6       if  $|\varphi_{\text{dec}}| > 0$  then
7            $Q_{\text{decision}}.\text{add}((id, |\varphi_{\text{dec}}| - 1, \varphi_{\text{dec}}.\text{last}()))$ 
8            $\langle F, \varphi_{\text{imp}} \rangle := \text{simplify\_and\_learn}(F, \varphi_{\text{dec}}, \varphi_{\text{imp}})$ 
9           if  $\varphi_{\text{dec}} \cup \varphi_{\text{imp}}$  falsify a clause in  $F$  then return UNSAT
10          if  $\varphi_{\text{dec}} \cup \varphi_{\text{imp}}$  assign all variables in  $F$  then return SAT
11           $l_{\text{dec}} := \text{decide}(F, \varphi_{\text{dec}}, \varphi_{\text{imp}})$ 
12          if  $LA\_search(F, \varphi_{\text{dec}} \cup \{l_{\text{dec}}\}, \varphi_{\text{imp}}, S_{\text{id}}) = \text{SAT}$  then return SAT
13          return  $LA\_search(F, \varphi_{\text{dec}} \cup \{\neg l_{\text{dec}}\}, \varphi_{\text{imp}}, S_{\text{id}})$ 

```

Fig. 2. Pseudocode listing of `march_ccc`.

4.3 Reintroducing the cutoff heuristic

One advantage of CC was that the conquer phase could be parallelized efficiently by using multiple CDCL solvers in parallel, each solving a single cube. With CCC_∞ this is no longer possible, since the lookahead solver will continue with a single branch until it is solved by either CDCL or lookahead. Additionally, CCC_∞ always uses twice as much CPU time as wall clock time, because the lookahead and CDCL solvers run in parallel.

To reduce this wasted resource utilization and allow for parallelization of the CDCL solver, we reintroduce the conquer phase by applying a suitable cutoff heuristic. As with CC, we pass cubes from the cube phase to the conquer phase via the file system using the iCNF⁶ format, which is basically a concatenation of the original formula F and the generated cubes as assumptions. An incremental SAT solver iterates over each cube c_{id} in the file, and solves $F \wedge c_{id}$ until a solution is found or all cubes have been refuted. We use iMiniSAT and iLingeling with four CDCL solvers for the serial respectively parallelized conquer phase, denoted CCC_{mini} and CCC_{lg14} . We use CCC to refer to the cube phase regardless of what conquer solver is used.

⁶ <http://users.ics.tkk.fi/swiering/icnf>

```

CDCL_search( $F$ )
1   $S := \emptyset$ 
2  forever do
3      if  $\varphi$  assigns all variables in  $F$  then return SAT
4      if  $Q_{\text{decision}} \neq \emptyset$  then
5           $\langle id, backtrackLevel, l_{\text{dec}} \rangle := Q_{\text{decision}}.\text{remove}()$ 
6          if  $backtrackLevel > |S|$  then continue
7           $S.\text{push}(id)$ 
8           $\varphi := \text{backtrack}(\varphi, backtrackLevel)$ 
9           $\varphi := \varphi \cup \{l_{\text{dec}} = \mathbf{t}\}$ 
10     else
11          $\varphi := \varphi \cup \{\text{pickDecisionLiteral}() = \mathbf{t}\}$ 
12      $\varphi := \text{propagate}(F, \varphi)$ 
13     while  $\varphi$  falsifies a clause in  $F$  do
14          $\langle conflict, backtrackLevel \rangle := \text{analyze}(F, \varphi)$ 
15         if  $conflict = \emptyset$  then return UNSAT
16         if  $backtrackLevel < |S|$  then
17             while  $|S| > backtrackLevel + 1$  do  $S.\text{pop}()$ 
18              $Q_{\text{solved}}.\text{add}(S.\text{pop}())$ 
19              $F := F \wedge conflict$ 
20              $\varphi := \text{backtrack}(\varphi, backtrackLevel)$ 
21              $\varphi := \text{propagate}(F, \varphi)$ 

```

Fig. 3. Pseudocode listing of minisat_ccc.

The cutoff heuristic of CC is based on a rough prediction of the performance of CDCL on a cube. Given a cube c_{id} , it computes its difficulty⁷⁸ $d(c_{id}) := |\varphi_{\text{dec}}|^2 \cdot (|\varphi_{\text{dec}}| + |\varphi_{\text{imp}}|)/n$, where $|\varphi_{\text{dec}}|$ and $|\varphi_{\text{imp}}|$ are the number of decision and implied variables respectively, and n is the total number of free variables. If $d(c_{id})$ is high, the CDCL solver is expected to solve c_{id} fast.

The cutoff heuristic in CC focuses on identifying cubes that are easy for CDCL to solve. It cuts off a branch if $d(c_{id})$ exceeds a dynamic threshold value t_{cc} . Initially $t_{cc} = 1000$, and it is multiplied by 0.7 whenever lookahead solves a cube (because it assumes that CDCL would have solved this cube faster) or when the number of decisions becomes too high (to avoid generating too many cubes). It is incremented by 5% at every decision to avoid the value from dropping too low.

For CCC, the same heuristic does not work because easy cubes are solved quickly by the CDCL solver. This makes the threshold very unstable so that it quickly converges to 0 or infinity depending on the instance. We therefore use a different heuristic, but using the same difficulty metric $d(c_{id})$.

Easy cubes can be detected better by CCC than by CC, because CCC can detect for which cubes CDCL finds a solution before the lookahead solver does. CCC would ideally cut off these cubes so that they can be solved in parallel. The contrary goes for when the lookahead solver solves a cube: it then seems

⁷ CC's heuristic has been improved slightly since it was initially published [4]; it now uses $|\varphi_{\text{dec}}|^2$ instead of $|\varphi_{\text{dec}}|$.

⁸ The notation is ours.




Fig. 4. Scatter plots showing the effect of the performance in seconds of MiniSAT (vertical axis) and CCC_{mini} (horizontal axis) on benchmarks selected (\triangle) and not selected ($+$) by the predictor. Left application benchmarks of SAT 2009 and SAT 2011, right crafted instances of SAT 2009 and SAT 2011. Above the line CCC_{mini} is stronger, below the line MiniSAT is stronger.

that lookahead contributes to the search, which means that it is not desirable to cut off.

CCC uses the same difficulty metric $d(c_{id})$ as CC, but a different heuristic for determining the threshold value t_{ccc} . If a cube c_{id} is solved by CDCL, the value is updated towards $s := 0.4 \cdot d(c_{id})$, whereas it is updated towards $s := 3 \cdot d(c_{id})$ if c_{id} was solved by lookahead. To avoid too sudden changes, t_{ccc} is not changed to s directly but is filtered by $t'_{ccc} := 0.4 \cdot s + 0.6 \cdot t_{ccc}$. To furthermore avoid the threshold from dropping too low, it is incremented for every cube that is cut off.

5 Empirical results

In this section we discuss the performance of the CCC solvers and the effectiveness prediction. We have first run CCC_{∞} for 5 seconds on all instances from the application and crafted categories of the SAT 2009 and 2011 competition⁹ and selected only instances where CCC_{∞} is not aborted in favor of a pure CDCL search by the predictor. These instances are referred to as *predicted* instances. Since the prediction takes at most 5 seconds and usually much less, we consider the overhead hardly significant. We therefore focus our experiments on predicted instances.¹⁰

The predictor selects 44 out of 292 instances from the SAT 2009 application suite, and 41 out of 300 from the SAT 2011 application suite. For crafted instances it selects a larger fraction: 70 out of 281 and 99 out of 276 for the 2009 and 2011 crafted suites respectively. As seen in Fig. 4, the predictor mostly selects instances for which CCC_{mini} works well compared to MiniSAT, and there

⁹ <http://www.satcompetition.org/>

¹⁰ The sources of the used software and the list of predicted instances are available on <http://fmv.jku.at/cccreview>.

are almost no instances where CCC_{mini} times out (>900 seconds) and MiniSAT does not. For unpredicted instances combined from both application categories, CCC_{mini} solves only 208 instances within a 900 second timeout versus 274 by MiniSAT. For the crafted instances that is 115 for CCC_{mini} versus 141 for MiniSAT. We therefore argue that the predictor is very well suited to select the instances where cube-and-conquer works well.

We ran each predicted instance on the following solvers: (C) CC_{mini} , (C) CC_{lg14} , CCC_{∞} , reference solvers MiniSAT 2.2, March_{rw} [8], and Lingeling, and parallel solver Plingeling4 (Plingeling with four threads). The CCC solvers all use MiniSAT 2.2 and March_{cc} (March_{rw} with cube support) concurrently in the cube phase as described in Sec.4, and the CC solvers only use March_{cc} in the cube phase. CCC_{mini} and CC_{mini} use iMiniSAT in the conquer phase, and (C) CC_{lg14} uses iLingeling with four parallel CDCL solvers in the conquer phase. Before passing an instance to any solver, the instance was preprocessed with Lingeling’s -s option.

We report on wall clock time unless stated otherwise. For CCC_{∞} the CPU time is twice the wall clock time since two solvers run concurrently. For CCC_{mini} the cube phase is usually short and the run time is dominated by the conquer phase, hence the wall clock and CPU time are often similar. For CCC_{lg14} the times deviate most, as the conquer phase is parallelized efficiently.

The cactus plots in Fig. 5 and Fig. 6 show that all cube-and-conquer techniques are strong on the predicted instances: all (C)CC solvers outperform the




Fig. 5. Cactus plot of various solvers on the application benchmarks of the SAT09 and SAT11 competitions selected by the CCC predictor.




Fig. 6. Cactus plot of various solvers on the crafted benchmarks of the SAT09 and SAT11 competitions selected by the CCC predictor.

three reference solvers in the crafted categories, and perform slightly better on application instances especially with lower timeout values. CCC_{mini} solves 3 more instances than MiniSAT within a 3600 second time limit for both application suites: 32 vs 35 and 25 vs 28 for the SAT 2009 and 2011 application instances respectively. The performance on crafted instances is even better: CCC_{mini} solves 5 more instances than MiniSAT (52 vs 57) in the SAT 2009 crafted category, and 17 more in the SAT 2011 crafted category (38 vs 55).

The results for CCC_{lgl4} show that the cubes generated by CCC can be parallelized well, even though CCC with a single Lingeling solver (not plotted for clarity) performs worse than CCC_{mini} . For application instances, the differences between Plingeling4 (the winner of the SAT 2011 competition’s application category in wall-clock time), CC_{lgl4} , and CCC_{lgl4} are not so large: CCC_{lgl4} performs slightly better for lower time limits, but Plingeling4 and CC_{lgl4} eventually solve one more instance. For crafted instances, CCC_{lgl4} performs best, and solves 7 more than its predecessor CC_{lgl4} and 24 more than Plingeling4.

It is interesting to see that CCC_{∞} still performs reasonably well, even though it is a very extreme version of CCC where the cube and conquer phases are fully merged. Although it is not the best configuration, it shows that the online usage of CDCL really contributes to the lookahead search: by cutting off leafs early using MiniSAT, CCC_{∞} solves many more instances than pure March does. Detailed results show that the wall clock time of CCC_{∞} is often slightly higher than CCC_{mini} , but the biggest problem is that some instances that are solved quickly by other solvers and not at all by CCC_{∞} . It seems that for some instances,

CDCL is not fast enough to cut off enough cubes. Additionally the CPU time is much larger for CCC_∞ because at all times the two solvers run concurrently without idling.

6 Conclusion

In this work we proposed an online cube-and-conquer solver that solves the two main limitations of offline cube-and-conquer. First, it is able to predict efficiently on which instances it works well, and abort the search after a few seconds in favor of a pure CDCL solver if not. Second, it does not estimate the performance of CDCL on a cube merely by assuming that it is similar to the performance of lookahead on that cube. This is not true in general so that offline cube-and-conquer is often not able to determine when to stop partitioning and start solving.

The cube-and-conquer solver we proposed runs a lookahead and CDCL solver concurrently to partition the search space. We have seen that this not only implicitly improves the run time of the cube phase, it also allows for better cutoff heuristics so that the generated cubes are easier for a CDCL solver to solve. Like offline cube-and-conquer, our approach allows the conquer phase to be parallelized efficiently.

References

1. N. Eén and N. Sörensson. Temporal induction by incremental SAT solving. *ENTCS*, 89(4):543–560, 2003.
2. N. Eén and N. Sörensson. An extensible SAT-solver. In *SAT’03*, volume 2919 of *LNCS*, pages 502–518. Springer, 2004.
3. M. J. H. Heule. *SmArT Solving: Tools and techniques for satisfiability solvers*. PhD thesis, Delft University of Technology, 2008.
4. M. J. H. Heule, O. Kullmann, S. Wieringa, and A. Biere. Cube and conquer: Guiding CDCL SAT solvers by lookaheads. In *Proc. HVC’11*, 2011. To be published.
5. M. J. H. Heule, J. E. van Zwieten, M. Dufour, and H. van Maaren. March_{eq}: Implementing additional reasoning into an efficient lookahead sat solver. In Holger H. Hoos and David G. Mitchell, editors, *SAT 2004*, volume 3542 of *Lecture Notes in Computer Science*, pages 345–359. Springer, 2005.
6. S. Kadioglu, Y. Malitsky, A. Sabharwal, H. Samulowitz, and M. Sellmann. Algorithm selection and scheduling. In Jimmy Ho-Man Lee, editor, *CP*, volume 6876 of *LNCS*, pages 454–469. Springer, 2011.
7. J. P. Marques-Silva, I. Lynce, and S. Malik. *Conflict-Driven Clause Learning SAT Solvers*, volume 185 of *FAIA*, chapter 4, pages 131–153. IOS Press, February 2009.
8. S. Mijnders, B. de Wilde, and M. J. H. Heule. Symbiosis of search and heuristics for random 3-sat. In David Mitchell and Eugenia Ternovska, editors, *Proceedings of the Third International Workshop on Logic and Search (LaSh 2010)*, 2010.
9. M. Nikolic, F. Maric, and P. Janicic. Instance-based selection of policies for sat solvers. In O. Kullmann, editor, *SAT*, volume 5584 of *LNCS*, pages 326–340. Springer, 2009.
10. L. Xu, F. Hutter, H. H. Hoos, and K. Leyton-Brown. Satzilla: Portfolio-based algorithm selection for sat. *J. Artif. Intell. Res. (JAIR)*, 32:565–606, 2008.